Mutual Information in Rank-One Matrix Estimation

Florent Krzakala*†§, Jiaming Xu§ and Lenka Zdeborová‡§
* Laboratoire de Physique Statistique, Sorbonnes Universités & Université Pierre & Marie Curie, 75005 Paris
† Laboratoire de Physique Statistique (CNRS UMR-8550), PSL Universités & École Normale Supérieure, 75005 Paris
§ Simons Institute for the Theory of Computing, University of California, Berkeley, Berkeley, CA, 94720
‡ Institut de Physique Théorique, CNRS, CEA, Université Paris-Saclay, F-91191, Gif-sur-Yvette, France.

arXiv:1603.08447v1 [cs.IT] 28 Mar 2016

***Abstract*—We consider the estimation of a n-dimensional vector $\boldsymbol{x}$ from the knowledge of noisy and possibility non-linear element-wise measurements of $\boldsymbol{x}\boldsymbol{x}^T$, a very generic problem that contains, e.g. stochastic 2-block model, submatrix localization or the spike perturbation of random matrices. We use an interpolation method proposed by Guerra [1] and later refined by Korada and Macris [2]. We prove that the Bethe mutual information (related to the Bethe free energy and conjectured to be exact by Lesieur et al. [3] on the basis of the non-rigorous cavity method) always yields an upper bound to the exact mutual information. We also provide a lower bound using a similar technique. For concreteness, we illustrate our findings on the sparse PCA problem, and observe that (a) our bounds match for a large region of parameters and (b) that it exists a phase transition in a region where the spectum remains uninformative. While we present only the case of rank-one symmetric matrix estimation, our proof technique is readily extendable to low-rank symmetric matrix or low-rank symmetric tensor estimation.**

I. Introduction and Main results

The estimation of low-rank matrices from their noisy, incomplete, or non linear measurements is a problem that has a wide range of applications of practical interest in machine learning and statistics, ranging from the sparse PCA [4], and community detection [5] to sub-matrix localization [6, 7]. We shall consider the setting where the rank-one matrix to be estimated is created as:

$$\boldsymbol{W} = \frac{1}{\sqrt{n}}\boldsymbol{x}\boldsymbol{x}^T\,, \tag{1}$$

where $\boldsymbol{x}$ is a n-dimensional vector whose elements were chosen independently at random from a prior distribution $p(x)$. The matrix $\boldsymbol{W}$ is then observed element-wise through a noisy non-linear output channel $P_{\rm out}(Y_{ij}|W_{ij})$, with $i,j=1,\dots,n$. We assume the noise to be symmetric so that $Y_{ij}=Y_{ji}$. The goal is to estimate the unknown vector $\boldsymbol{x}$ from $\boldsymbol{Y}$ up to a global flip of sign, or equivalently the unknown rank-one matrix $\boldsymbol{W}$ from $\boldsymbol{Y}$. Throughout this paper, we assume that $p(x)$ and $P_{\rm out}(y,w)$ are independent of n, and are known.

We consider here an information-theoretic viewpoint and analyze the mutual information for the above model defined as $I(\boldsymbol{W};\boldsymbol{Y}) = I(\boldsymbol{x};\boldsymbol{Y}) = \mathbb{E}_{\boldsymbol{x},\boldsymbol{Y}}\log\{P(\boldsymbol{x},\boldsymbol{Y})/[p(\boldsymbol{x})P(\boldsymbol{Y})]\}$. Up to a simple term, see eq. (10), the mutual information is related to the free energy, which is the fundamental quantity usually considered in statistical physics [8, 9]. Recently, an explicit single-letter characterization of the mutual information between the noisy observation and the vector to be recovered has been computed in some special cases of our setting [4, 5]. The general formula has been derived by Lesieur *et al.* [3, 10] on the basis of the heuristic cavity method from statistical mechanics [8, 9, 11]. We shall refer to the formula conjectured in [3] as the *Bethe mutual information*. In this contribution, we use a rigorous technique that also originated in physics, the so-called Guerra interpolation [1, 2], to prove that the Bethe mutual information provides always an upper bound. By a variant of the Guerra interpolation, we also provide a lower bound on the mutual information that matches the upper bound for a sizable range of parameters.

A. Main results

Our first result is a rigorous proof of a conjecture from Lesieur *et al.* [3] of channel universality. In the context of community detection in graphs with growing average degrees, an equivalence between Bernoulli channel and Gaussian channel has been proven already in [5].

Theorem I.1 (Channel Universality)**.** *Assume model (1) with a prior $p(x)$ having a finite support, and the output channel $P_{\rm out}(y|w)$ such that at $w=0$, $\log P_{\rm out}(y|w)$ is thrice differentiable with bounded second and third derivatives and $\mathbb{E}_{P_{\rm out}(y|0)}[|\partial_w \log P_{\rm out}(y|w)|_{w=0}|^3] = O(1)$. Then the mutual information per variable satisfies*

$$I(\boldsymbol{W};\boldsymbol{Y}) = I(\boldsymbol{W};\boldsymbol{W}+\sqrt{\Delta}\,\boldsymbol{\xi}) + O(\sqrt{n}), \tag{2}$$

where $\boldsymbol{\xi}$ is a symmetric matrix such that $\xi_{ij} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$ for $i\le j$, and Δ is the inverse Fisher information (evaluated at $w=0$) of the channel $P_{\rm out}(y|w)$:

$$\frac{1}{\Delta} \equiv \mathbb{E}_{P_{\rm out}(y|0)}\left[\left(\frac{\partial \log P_{\rm out}(y|w)}{\partial w}\Big|_{y,0}\right)^2\right]. \tag{3}$$

Informally, this means that we only have to compute the mutual information for a Gaussian channel to take care of a wide range of channels.

Our next result is that, the Bethe mutual information is always an upper bound to the true one for any finite n:

Theorem I.2 (Upper Bound)**.** *Assume model (1) with a prior $p(x)$ having finite support, and a Gaussian channel such that $P_{\rm out}(y|w)$ is the probability density function of a centered Gaussian distribution with variance Δ. Then for all non-negative parameter m, the mutual information per variable $I(\boldsymbol{x};\boldsymbol{Y})/n$ is upper bounded by the Bethe mutual information*

$i_{\rm B}(m)$ *defined by*

$$i_{\rm B}(m)=\frac{m^2+\left[\mathbb{E}_x(x^2)\right]^2}{4\Delta}-\mathbb{E}_{x,z}\left[\mathcal{J}\left(\frac{m}{\Delta},\frac{mx}{\Delta}+\sqrt{\frac{m}{\Delta}}z\right)\right], \quad (4)$$

where $\mathbb{E}_x$ *denotes the expectation taken over the prior distribution* $p(x)$*,* z *is a Gaussian variable following* $\mathcal{N}(0,1)$*, and where*

$$\mathcal{J}(A,B)=\log\int e^{Bx-Ax^2/2}p(x)\mathrm{d}x\,. \quad (5)$$

Note that the mutual information is related to the free energy only by a simple term eq. (10).

Our last main result yields an asymptotic lower bound:

Theorem I.3 (Lower Bound)**.** *With the same hypothesis as in Theorem I.2, denote* $\widehat{m}$ *the minimizer of eq. (4). Define*

$$i_{\rm L}(m)=\frac{2m^2-\widehat{m}^2+\left[\mathbb{E}_x(x^2)\right]^2}{4\Delta}-\mathbb{E}_{x,z}\left[\mathcal{J}\left(\frac{\widehat{m}}{\Delta},\frac{mx}{\Delta}+\sqrt{\frac{\widehat{m}}{\Delta}}z\right)\right], \quad (6)$$

where $z\sim\mathcal{N}(0,1)$*. Assume that* $i'_{\rm L}(m)=0$ *has a finite number of solutions. Then*

$$\liminf_{n\to\infty}\frac{1}{n}I(\boldsymbol{x};\boldsymbol{Y})\geq\min_m i_{\rm L}(m).$$

One can verify that $\widehat{m}$ is always a stationary point of $i_{\rm L}(m)$. If additionally $\widetilde{m}=\widehat{m}$, where $\widetilde{m}\equiv\operatorname{argmin}\, i_{\rm L}(m)$, then $\min_m i_{\rm L}(m)=i_{\rm B}(\widehat{m})$ and the Bethe mutual information asymptotically equals the true one. As we shall see this is the case for some range of parameters, but not always.

B. Relation to previous works

For two particular cases of model (1), the mutual information has previously been proven rigorously [4, 5] using the approximate message passing algorithm and its state evolution [12]. Remarkably, these were constructive proofs, with an explicit algorithm that achieves the minimum mean squared error (MMSE). The proof technique of [4, 5] does not extend straightforwardly when the state evolution had more than one fixed point. Our approach applies to more general class of problems (even when several fixed points exist), but is not constructive and our lower bound is not always tight.

We rely on two essential contributions. First, we use the cavity computations of Lesieur *et al.* who solved the problem using statistical physics methods [3]. Our results are a considerable step towards confirming the full validity of this approach. We show that the Bethe mutual information always yields an upper bound, and by our lower bound we confirm that the Bethe mutual information is exact for a large range of parameters. Secondly, our approach is inspired by the scheme introduced by Korada and Macris [2] for studying the Sherrington-Kirkpatrick model of spin glasses on the so-called Nishimori line with the Guerra interpolation [1]. It also crucially exploits the Nishimori identities [8] for optimal Bayesian estimation [13–15]. It is worth to remark that for problems of Bayes-optimal estimation (on the Nishimori line) the simplest version of the Guerra interpolation provides an upper bound on the free energy/mutual information while

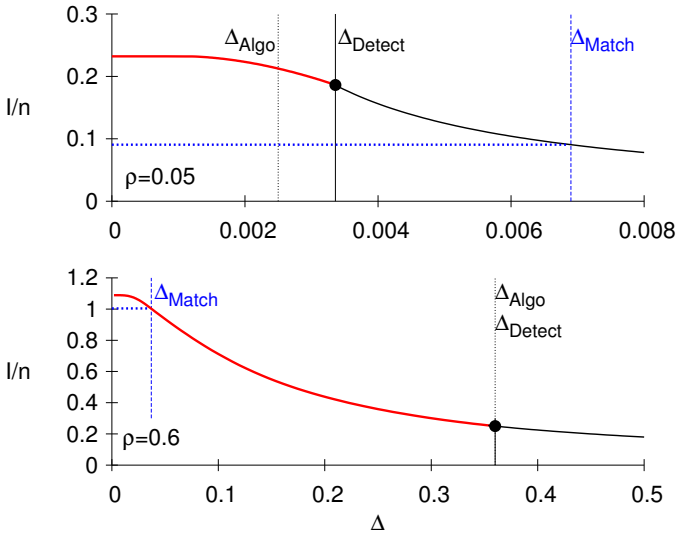

Fig. 1. Asymptotic mutual information I/n per variable as a function of the effective channel noise Δ for sparse Rademacher variables with density $\rho=0.05$ (top) and $\rho=0.6$ (bottom). Both bounds are tight and equal to the Bethe mutual information for $\Delta>\Delta_{\rm Match}(\rho)$. For $\Delta<\Delta_{\rm Match}$, however, the lower bound (shown in blue dashed line) is not tight and does not coincide with the (conjectured exact) upper bound (full line). The detectability transition in the upper bound arises at $\Delta_{\rm Detect}$, and is a (conjectured tight) lower bound on the true detectability transition. For large enough ρ, the bounds are tight at the phase transition and we further observe that $\Delta_{\rm Detect}=\Delta_{\rm Algo}$ (for instance here when $\rho=0.6$). However, we find that when $\rho\approx 0.09$, $\Delta_{\rm Detect}$ becomes strictly larger than $\Delta_{\rm Algo}$, thus revealing the existence of a regime where detection is informationally possible but computationally hard for known polynomial-time algorithms.

for standard statistical mechanics models, or for optimization problems, it provides instead a lower bound [1, 16].

While we present only the rank-one version on model (1), our proof is readily extendable to any finite rank, or even to tensor factorization. Future directions include the extension to non-symmetric matrices, which are less straightforward. We believe that our results, once more, give strong credibility to the use of the replica and the cavity methods for statistical estimation problems.

II. Application to sparse Rademacher variables

We shall illustrate (see Fig. 1), for concreteness, our findings on a specific example of sparse Rademacher variables where $x=0,+1,-1$ with probability $1-\rho,\rho/2,\rho/2$ respectively.

Solving numerically for (4) and (6) shows that the region where the two bounds coincide, and the Bethe mutual information is thus rigorously exact, is quite sizable. This happens, for instance, for all ρ if Δ is larger than a value around 0.15, and for all Δ if ρ is larger than a value around 0.66. For large enough Δ, we also find that $\widehat{m}=\widetilde{m}=0$, and a derivative of the mutual information indicates that it is impossible to find an assignment correlated to the truth.

In the non-sparse case ($\rho=1$), one can further show that the minimizers of both (4) and (6) always coincide and the Bethe mutual information is thus asymptotically exact for all $\Delta>0$. A phase transition arises at $\Delta=1$ so that for $\Delta<1$, a non trivial solution with $\widehat{m}=\widetilde{m}>0$ appears. Among the problems

that belong into this category are the dense version of the binary stochastic block model [5, 17, 18], the dense version of the censored block model [19–21], or the Sherrington-Kirkpatrick model on the Nishimori line as originally studied by Korada and Macris [2].

When both ρ and Δ are small enough, numerically we find that our upper and lower bounds stop to coincide. Define $\Delta_{\rm match}(\rho)$ as the minimum Δ so that our upper and lower bounds match, for a fixed ρ. This is illustrated in Fig. 1 with three different values of ρ. When $\Delta < \Delta_{\rm Algo} = \sqrt{\rho}$, polynomial-time algorithms such as message passing [3, 5, 10, 12] or spectral methods [22] are known to be able to find an assignment with a non-trivial correlation to the truth; thus, in this region, the non-trivial detection is easy. We define $\Delta_{\rm Detect}$ to be the (conjectured) information-theoretic (IT) threshold for the non-trivial detection. Depending on the particular values of ρ, we have the following two strikingly different observations:

(1) When ρ is large (e.g. $\rho = 0.6$), $\Delta_{\rm Match} < \Delta_{\rm Algo} = \Delta_{\rm Detect}$. In this case, our lower and upper bound coincide, showing that there is a non-analyticity (phase transition) at $\Delta_{\rm Detect}$. In this region, the conjectured IT threshold is indeed the true one and coincides with the algorithmic one $\Delta_{\rm Algo}$.

(2) When ρ becomes smaller than a certain threshold ρ^*, $\Delta_{\rm Algo} < \Delta_{\rm Detect} < \Delta_{\rm Match}$. Numerically we find that $\rho^* \approx 0.09$. The derivative of $i_{\rm B}$ undergoes a phase transition at $\Delta_{\rm Detect}$. It readily implies that the derivative of the true mutual information per variable must exhibit a phase transition somewhere between $\Delta_{\rm Match}$ and $\Delta_{\rm Detect}$, which is strictly above $\Delta_{\rm Algo}$. Hence, in this case, there exists a region where the non-trivial detection is informationally possible, but it is impossible via standard polynomial-time algorithms like spectral methods or message passing.

III. Channel universality

Let us now show that in order to characterize the mutual information per variable, it suffices to consider an equivalent Gaussian channel. We give a detailed rigorous proof in Appendix A and present here only its main idea. We assume that the prior $p(x)$ has a finite support and denote

$$S_{ij} \equiv \partial_w \log P_{\rm out}(Y_{ij}|w)|_{w=0}\,, \tag{7}$$
$$S'_{ij} \equiv \partial_w^2 \log P_{\rm out}(Y_{ij}|w)|_{w=0}\,. \tag{8}$$

We assume $\mathbb{E}_{Y_{ij}|0}[|S_{ij}|^3]$, S'_{ij}, and $\partial_w^3 \log P_{\rm out}(Y_{ij}|w)|_{w=0}$ are all bounded.

Note that $W_{ij} = O(1/\sqrt{n})$. Using Taylor's expansion of $\log P_{\rm out}(y|w)|$ at $w = 0$, for all $i \le j$ we can write

$$P_{\rm out}(Y_{ij}|W_{ij}) = P_{\rm out}(Y_{ij}|0) e^{W_{ij}S_{ij} + \frac{1}{2}W_{ij}^2 S'_{ij} + O(n^{-3/2})}.$$

Thus,

$$P_{\rm out}(\boldsymbol{Y}|\boldsymbol{W}) = P_{\rm out}(\boldsymbol{Y}|0) e^{\sum_{i\le j}(W_{ij}S_{ij} + \frac{1}{2}W_{ij}^2 S'_{ij}) + O(\sqrt{n})}, \tag{9}$$

Classical properties of the Fisher information give that $\mathbb{E}_{Y_{ij}|0}[S'_{ij}] = -\mathbb{E}_{Y_{ij}|0}[S_{ij}^2] = -1/\Delta$. Using the fact that $P_{\rm out}(y|w)$ is close to $P_{\rm out}(y|0)$, one can further argue that $\mathbb{E}_{Y_{ij}|W_{ij}}[S'_{ij}] = -1/\Delta + O(n^{-1/2})$. By concentration inequalities, we expect that

$$\sum_{i\le j} W_{ij}^2 S'_{ij} \approx \sum_{i\le j} W_{ij}^2 \mathbb{E}_{Y_{ij}|W_{ij}}[S'_{ij}] = \sum_{i\le j} W_{ij}^2/\Delta + O(\sqrt{n}).$$

Thus

$$P_{\rm out}(\boldsymbol{Y}|\boldsymbol{W}) \approx P_{\rm out}(\boldsymbol{Y}|0) e^{\sum_{i\le j}(W_{ij}S_{ij} - \frac{1}{2\Delta}W_{ij}^2) + O(\sqrt{n})}$$

and consequently

$$P(\boldsymbol{W}|\boldsymbol{Y}) \propto P(\boldsymbol{W}) e^{-\frac{1}{2\Delta}\sum_{i\le j}(\Delta S_{ij} - W_{ij})^2 + O(\sqrt{n})}.$$

Hence, we expect that

$$I(\boldsymbol{W};\boldsymbol{Y}) = I(\boldsymbol{W};\boldsymbol{W} + \sqrt{\Delta}\,\boldsymbol{\xi}) + O(\sqrt{n}).$$

In other words, the mutual information per variable $I(\boldsymbol{x};\boldsymbol{Y})/n$ is asymptotically equal to the mutual information per variable of a Gaussian channel with noise variance Δ.

IV. Proving the upper bound

A. Mutual information and free energy

Using the channel universality, we only need to deal with the Gaussian output. The mutual information between the observation $\boldsymbol{Y}$ and the unknown vector $\boldsymbol{x}$ is defined using the entropy as $I(\boldsymbol{x};\boldsymbol{Y}) = H(\boldsymbol{Y}) - H(\boldsymbol{Y}|\boldsymbol{x})$. For the Gaussian noise, a straightforward computation shows that the mutual information per variable is expressed as

$$\frac{I(\boldsymbol{x};\boldsymbol{Y})}{n} = f + \frac{[\mathbb{E}_x(x^2)]^2}{4\Delta}\,, \tag{10}$$

where $f = -\mathbb{E}_{\boldsymbol{Y}}\left[\log Z(\boldsymbol{Y})\right]/n$ is the average free energy per variable using the standard statistical physics terms, and $Z(\boldsymbol{Y})$ is the partition function defined by

$$Z(\boldsymbol{Y}) \equiv \int \mathrm{d}\boldsymbol{x}\, p(\boldsymbol{x}) \exp\left[\sum_{i\le j}\left(-\frac{x_i^2 x_j^2}{2n\Delta} + \frac{x_i x_j Y_{ij}}{\sqrt{n}\Delta}\right)\right]. \tag{11}$$

We now show how to upper bound the free energy f.

B. Denoising

We first solve a simpler denoising problem. Assume we observe a noisy version of a vector $\boldsymbol{x}^*$ that we denote $\boldsymbol{y}$:

$$\boldsymbol{y} \;=\; \boldsymbol{x}^* + \sigma \boldsymbol{z}, \tag{12}$$

where $x_i^* \overset{\text{i.i.d.}}{\sim} p(x)$ and $z_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$. The corresponding posterior distribution reads

$$P(\boldsymbol{x}|\boldsymbol{y}) = \frac{1}{Z_0} p(\boldsymbol{x}) \exp\left(-\frac{\|\boldsymbol{x}\|_2^2}{2\sigma^2} + \frac{\langle \boldsymbol{x}, \boldsymbol{y}\rangle}{\sigma^2}\right). \tag{13}$$

For future convenience we denote the variance $\sigma^2 \equiv \frac{\Delta}{m}$, where m and Δ are some so-far unspecified parameters. For this denoising problem, the averaged free energy per variables reads

$$-n f_0 = \mathbb{E}[\log Z_0] = \mathbb{E}_{x^*,z}\left[\mathcal{J}\left(\frac{m}{\Delta}, \frac{m x^*}{\Delta} + \sqrt{\frac{m}{\Delta}} z\right)\right], \tag{14}$$

where $x^* \sim p(x)$, $z \sim \mathcal{N}(0,1)$, and $\mathcal{J}(A,B)$ is the function defined in eq. (5). Notice how this yields a formula very close to the one in Theorem I.3.

C. The interpolation method

We now use the Guerra interpolation method, setting an artificial parameter t, where we interpolate between the denoising problem at $t=0$ and the desired matrix factorization one at $t=1$. To do so, assume that we have access to two types of noisy observations: (1) A noisy version of $\boldsymbol{x}^*$, as in eq. (12), with now $\sigma^2 = \frac{\Delta}{m(1-t)}$; and (2) a noisy version of $W_{ij} = x_i^* x_j^*/\sqrt{n}$ with a Gaussian noise of variance Δ/t. The posterior distribution, in this case, is given by

$$P_t(\boldsymbol{x}|\boldsymbol{Y},\boldsymbol{y}) = \frac{1}{Z_t} p(\boldsymbol{x}) e^{t\sum_{i\le j}\left[-\frac{x_i^2 x_j^2}{2n\Delta}+\frac{x_i x_j Y_{ij}}{\sqrt{n}\Delta}\right]+(1-t)\left[-\frac{m\|\boldsymbol{x}\|_2^2}{2\Delta}+\frac{m\langle \boldsymbol{x},\boldsymbol{y}\rangle}{\Delta}\right]}.$$

This model interpolates between the denoising problem at $t=0$ and the one of the matrix factorization problem at $t=1$. Using the fundamental theorem of algebra, we write

$$-nf_1 = \mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\log Z_1\right] = f_0 - \int_0^1 dt \frac{d}{dt}\mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\log Z_t\right]. \tag{15}$$

The free energy at $t=0$ is precisely given by (14). Using now eq. (10) and eq. (4) we write

$$\frac{I(\boldsymbol{x};\boldsymbol{Y})}{n} = i_{\rm B}(m) - \frac{m^2}{4\Delta} - \frac{1}{n}\int_0^1 dt \frac{d}{dt}\mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\log Z_t\right]. \tag{16}$$

Theorem I.2 follows from the following lemma:

Lemma IV.1. *For all positive n and $t\in[0,1]$, we have*

$$\frac{1}{n}\frac{d}{dt}\mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\log Z_t\right] \ge -\frac{m^2}{4\Delta}. \tag{17}$$

D. The proof

Define

$$\mathcal{H}_t(\boldsymbol{x},\boldsymbol{x}^*,\boldsymbol{\xi},\boldsymbol{z}) = \sum_{i\le j} -\frac{t x_i^2 x_j^2}{2\Delta n} + \frac{t x_i x_j x_i^* x_j^*}{\Delta n} + x_i x_j\sqrt{\frac{t}{n\Delta}}\xi_{ij} - \frac{(1-t)m\|\boldsymbol{x}\|_2^2}{2\Delta} + \frac{(1-t)m\langle \boldsymbol{x},\boldsymbol{x}^*\rangle}{\Delta} + \langle \boldsymbol{x},\boldsymbol{z}\rangle\sqrt{\frac{m(1-t)}{\Delta}}.$$

Then $P_t(\boldsymbol{x}|\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}) = p(\boldsymbol{x})e^{\mathcal{H}_t}/Z_t$. Now we need to compute $\frac{d}{dt}\mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\log Z_t\right]$. Notice that

$$\partial_t \mathcal{H}_t = \sum_{i\le j}\mathcal{A}_{ij} + \sum_i \mathcal{B}_i.$$

where

$$\mathcal{A}_{ij} = -\frac{x_i^2 x_j^2}{2\Delta n} + \frac{x_i x_j x_i^* x_j^*}{\Delta n} + \frac{x_i x_j}{2\sqrt{n\Delta t}}\xi_{ij}\,, \tag{18}$$

$$\mathcal{B}_i = m\frac{x_i^2}{2\Delta} - m\frac{x_i x_i^*}{\Delta} - \frac{x_i z_i}{2}\sqrt{\frac{m}{\Delta(1-t)}}. \tag{19}$$

Since the prior $p(x)$ has a finite support $|\mathcal{A}_{ij}|$ and $|\mathcal{B}_i|$ are dominated by functions integrable with respect to P_t. Thus by the dominated convergence theorem,

$$\frac{d}{dt}\log Z_t = \mathbb{E}_t[\partial_t\mathcal{H}_t] = \sum_{i\le j}\mathbb{E}_t[\mathcal{A}_{ij}] + \sum_i \mathbb{E}_t[\mathcal{B}_i].$$

Moreover, $\mathbb{E}_t[\mathcal{A}_{ij}]$ and $\mathbb{E}_t[\mathcal{B}_i]$ are dominated by functions integrable with respect to the distribution of $\{\boldsymbol{x}^*,\boldsymbol{\xi},\boldsymbol{z}\}$. Thus again by the dominated convergence theorem,

$$\frac{d}{dt}\mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\log Z_t\right] = \mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\frac{d}{dt}\log Z_t\right] = \mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\sum_{i\le j}\mathbb{E}_t[\mathcal{A}_{ij}] + \sum_i \mathbb{E}_t[\mathcal{B}_i]\right]. \tag{20}$$

We then compute $\mathbb{E}_{\xi_{ij}}\left[\mathbb{E}_t[x_i x_j]\xi_{ij}\right]$ and $\mathbb{E}_{z_i}[\mathbb{E}_t[x_i]z_i]$. We use the integration by part to get rid of the z_i and ξ_{ij}. In particular, for a standard Gaussian random variable z and a continuous differentiable function $f(a)$ such that $f(a)\mathrm{e}^{-a^2/2}\to 0$ as $a\to\infty$, we have that $\mathbb{E}\left[zf(z)\right] = \mathbb{E}\left[f'(z)\right]$. Notice that P_t is a function of $\boldsymbol{\xi}$ and $\boldsymbol{z}$. Also, $\partial_{\xi_{ij}}\mathcal{H}_t = \sqrt{\frac{t}{n\Delta}}x_i x_j$. Then $|\partial_{\xi_{ij}}\mathcal{H}_t|$ is dominated by a function integrable under P_t. By the dominated convergence theorem,

$$\partial_{\xi_{ij}} Z_t = \sqrt{\frac{t}{n\Delta}} Z_t\ \mathbb{E}_t[x_i x_j].$$

It follows that

$$\partial_{\xi_{ij}}\mathbb{E}_t[x_i x_j] = \sqrt{\frac{t}{n\Delta}}\left(\mathbb{E}_t[x_i^2 x_j^2] - \left(\mathbb{E}_t[x_i x_j]\right)^2\right).$$

Thus $\partial_{\xi_{ij}}\mathbb{E}_t[x_i x_j]$ is continuous in ξ_{ij}. Applying the integration by parts, it yields that

$$\mathbb{E}_{\xi_{ij}}\left[\mathbb{E}_t[x_i x_j]\xi_{ij}\right] = \mathbb{E}_{\xi_{ij}}\left[\partial_{\xi_{ij}}\mathbb{E}_t[x_i x_j]\right] = \sqrt{\frac{t}{n\Delta}}\mathbb{E}_{\xi_{ij}}\left[\mathbb{E}_t[x_i^2 x_j^2] - \left(\mathbb{E}_t[x_i x_j]\right)^2\right].$$

Similarly, one can show that

$$\mathbb{E}_{z_i}[\mathbb{E}_t[x_i]z_i] = \mathbb{E}_{z_i}\left[\partial_{z_i}\mathbb{E}_t[x_i]\right] = \sqrt{\frac{m(1-t)}{\Delta}}\mathbb{E}_{z_i}\left[\mathbb{E}_t[x_i^2] - \left(\mathbb{E}_t[x_i]\right)^2\right].$$

It follows from (18) and (19) that

$$\mathbb{E}_{t,\boldsymbol{\xi}}[\mathcal{A}_{ij}] = \mathbb{E}_{t,\boldsymbol{\xi}}\left[\frac{x_i x_j x_i^* x_j^*}{\Delta n} - \frac{x_i x_j \mathbb{E}_t[x_i x_j]}{2n\Delta}\right], \tag{21}$$

$$\mathbb{E}_{t,\mathbf{z}}[\mathcal{B}_i] = \mathbb{E}_{t,\mathbf{z}}\left[-\tfrac{m}{\Delta}x_i x_i^* + \tfrac{m}{2\Delta}x_i\mathbb{E}_t[x_i]\right]. \tag{22}$$

Using the Nishimori identities given by Lemma A.4, we have that

$$\mathbb{E}_{\boldsymbol{x}^*,\mathbf{z},\boldsymbol{\xi}}\left[\mathbb{E}_t\left[x_i x_j\right]^2\right] = \mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\mathbb{E}_t\left[x_i x_j x_i^* x_j^*\right]\right], \tag{23}$$

$$\mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\left(\mathbb{E}_t\left[x_i\right]\right)^2\right] = \mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\mathbb{E}_t\left[x_i x_i^*\right]\right]. \tag{24}$$

Combining (20), (21)–(22), and (23)–(24) yields that

$$\begin{aligned}
&\frac{1}{n}\frac{d}{dt}\mathbb{E}_{\boldsymbol{x}^*,\boldsymbol{z},\boldsymbol{\xi}}\left[\log Z_t\right] \\
&= \frac{1}{2\Delta n^2}\sum_{i\le j}\mathbb{E}\left[x_i x_j x_i^* x_j^*\right] - \frac{m}{2\Delta n}\sum_i \mathbb{E}\left[x_i x_i^*\right] \\
&\ge \frac{\mathbb{E}[m_t^2]}{4\Delta} - \frac{m\mathbb{E}[m_t]}{2\Delta} \\
&= \frac{1}{4\Delta}\mathbb{E}[(m_t - m)^2] - \frac{m^2}{4\Delta} \ge -\frac{m^2}{4\Delta},
\end{aligned}$$

where $m_t = \langle \boldsymbol{x},\boldsymbol{x}^*\rangle$ with $\boldsymbol{x}$ drawn from P_t. Lemma IV.1 readily follows.

V. Proving the lower bound

The proof for the lower bound also relies on the interpolation method. Again, the proof idea is inspired by [2].

A. An ad-hoc model

We shall first compute the free energy of a totally artificial model, that does not correspond to any Bayesian inference problem. Later, we will interpolate the desired free energy starting from this model at $t = 0$. Let $\widehat{m}$ denote the minimizer of eq. (4). For a fixed set of $(\boldsymbol{z}, \boldsymbol{x}^*)$, let

$$\begin{aligned}\widetilde{Z}_0 &= \int \mathrm{d}\boldsymbol{x}\, p(\boldsymbol{x}) e^{\frac{1}{n\Delta}\sum_{i\le j}\left[x_i x_j x_i^* x_j^*\right] - \frac{\widehat{m}\|\boldsymbol{x}\|_2^2}{2\Delta} + \sqrt{\frac{\widehat{m}}{\Delta}}\langle \boldsymbol{x},\boldsymbol{z}\rangle} \\ &= \int \mathrm{d}\boldsymbol{x}\, p(\boldsymbol{x}) e^{\frac{1}{2n\Delta}\langle \boldsymbol{x},\boldsymbol{x}^*\rangle^2 - \frac{\widehat{m}}{2\Delta}\|\boldsymbol{x}\|_2^2 + \sqrt{\frac{\widehat{m}}{\Delta}}\langle \boldsymbol{x},\boldsymbol{z}\rangle + O(1)},\end{aligned}$$

where $O(1) = \frac{1}{2n\Delta}\|\boldsymbol{x}\|_2^2$.

Using the Gaussian identity $e^{\frac{b^2}{4a}} = \sqrt{a/\pi}\int e^{-am^2+bm}\mathrm{d}m$, with $a = n/2\Delta$ and $b = \langle \boldsymbol{x}, \boldsymbol{x^*}\rangle/\Delta$, we reach

$$\widetilde{Z}_0 \lesssim \sqrt{\frac{n}{2\pi\Delta}} \int \mathrm{d}\boldsymbol{x}\, p(\boldsymbol{x}) \int \mathrm{d}m e^{-\frac{nm^2}{2\Delta} + \langle \boldsymbol{x}, \frac{m}{\Delta}\boldsymbol{x^*} + \sqrt{\frac{\widehat{m}}{\Delta}}\boldsymbol{z}\rangle - \frac{\widehat{m}}{2\Delta}\|\boldsymbol{x}\|_2^2},$$

where $a_n \lesssim b_n$ means $a_n = O(b_n)$. We now invert the integral by Fubini's theorem so that

$$\widetilde{Z}_0 \lesssim \sqrt{\frac{n}{2\pi\Delta}} \int \mathrm{d}m \exp\left[\frac{-nm^2}{2\Delta} + \sum_i \mathcal{J}\left(\frac{\widehat{m}}{\Delta}, \frac{mx_i^*}{\Delta} + \sqrt{\frac{\widehat{m}}{\Delta}}z_i\right)\right].$$

Then, a naive application of the Laplace method suggests that

$$\widetilde{f}_0 := -\limsup_{n\to\infty} \frac{1}{n}\mathbb{E}_{\boldsymbol{z},\boldsymbol{x^*}}\left[\log \widetilde{Z}_0\right] \tag{25}$$

$$\ge \min_m \left[\frac{m^2}{2\Delta} - \mathbb{E}_{x^*,z}\left[\mathcal{J}\left(\frac{\widehat{m}}{\Delta}, \frac{mx^*}{\Delta} + \sqrt{\frac{\widehat{m}}{\Delta}}z\right)\right]\right]. \tag{26}$$

A rigorous proof that $\widetilde{f}_0$ is indeed lower bounded by the above expression (which is only what we require) is presented in Appendix C, under the assumption that $i'_{\mathrm{L}}(m) = 0$ has a finite number of solutions.

B. Interpolation reloaded

The proof of the lower bound then proceed again via the interpolation method, where we interpolate between the ad-hoc model and the matrix factorization one by considering the following partition function, at fixed value of $\{\boldsymbol{z}, \boldsymbol{\xi}, \boldsymbol{x^*}\}$:

$$\begin{aligned}\widetilde{Z}_t \;=\; &\int \mathrm{d}\boldsymbol{x}\, p(\boldsymbol{x}) e^{\frac{1}{n\Delta}\sum_{i\le j}\left[-t\frac{x_i^2 x_j^2}{2} + x_i x_j x_0^i x_0^j + \sqrt{n\Delta t}x_i x_j \xi_{ij}\right]} \\ &\times e^{(1-t)\left[-\frac{\widehat{m}\|\boldsymbol{x}\|_2^2}{2\Delta} + \left(\sqrt{\frac{\widehat{m}}{\Delta(1-t)}}\langle \boldsymbol{x},\boldsymbol{z}\rangle\right)\right]}.\end{aligned} \tag{27}$$

This is again detailed in Appendix D.

Acknowledgments

We thank J. Barbier, T. Lesieur, N. Macris and C. Moore for helpful discussions. Part of the research has received funding from the European Research Council under the European Union's 7th Framework Programme (FP/2007-2013/ERC Grant Agreement 307087-SPARCS).

References

[1] F. Guerra, "Broken replica symmetry bounds in the mean field spin glass model," *Communications in mathematical physics*, vol. 233, p. 1, 2003.

[2] S. B. Korada and N. Macris, "Exact solution of the gauge symmetric p-spin glass model on a complete graph," *Journal of Statistical Physics*, vol. 136, no. 2, pp. 205–230, 2009.

[3] T. Lesieur, F. Krzakala, and L. Zdeborová, "Mmse of probabilistic low-rank matrix estimation: Universality with respect to the output channel," *arXiv preprint arXiv:1507.03857*, 2015.

[4] Y. Deshpande and A. Montanari, "Information-theoretically optimal sparse pca," in *Information Theory (ISIT), 2014 IEEE International Symposium on*, June 2014, pp. 2197–2201.

[5] Y. Deshpande, E. Abbe, and A. Montanari, "Asymptotic mutual information for the two-groups stochastic block model," *arXiv:1507.08685*.

[6] Y. Chen and J. Xu, in *Proceedings of ICML 2014 (Also arXiv:1402.1267)*.

[7] B. Hajek, Y. Wu, and J. Xu, "Submatrix localization via message passing," *arXiv preprint arXiv:1510.09219*, 2015.

[8] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction*. Oxford University Press, 2001.

[9] M. Mézard and A. Montanari, *Information, Physics, and Computation*. Oxford: Oxford Press, 2009.

[10] T. Lesieur, F. Krzakala, and L. Zdeborová, "Phase transitions in sparse pca," in *Information Theory (ISIT), 2015 IEEE International Symposium on*. IEEE, 2015, pp. 1635–1639.

[11] M. Mézard, G. Parisi, and M. A. Virasoro. Singapore: World Scientific.

[12] S. Rangan and A. K. Fletcher, "Iterative estimation of constrained rank-one matrices in noise," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 1246–1250.

[13] Y. Iba, "The nishimori line and bayesian statistics," *Journal of Physics A: Mathematical and General*, vol. 32, no. 21, p. 3875, 1999.

[14] S. B. Korada and N. Macris, "Tight bounds on the capacity of binary input random cdma systems," *Information Theory, IEEE Transactions on*, vol. 56, no. 11, pp. 5590–5613, 2010.

[15] L. Zdeborová and F. Krzakala, "Statistical physics of inference: Thresholds and algorithms," *arXiv preprint arXiv:1511.02476*, 2015.

[16] S. Franz, M. Leone, and F. L. Toninelli, "Replica bounds for diluted non-poissonian spin systems," *Journal of Physics A: Mathematical and General*, vol. 36, no. 43, p. 10967, 2003.

[17] L. Massoulié, "Community detection thresholds and the weak ramanujan property," in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. ACM, 2014, pp. 694–703.

[18] E. Mossel, J. Neeman, and A. Sly, "A proof of the block model threshold conjecture," *arxiv:1311.4115*, 2013.

[19] E. Abbe, A. S. Bandeira, A. Bracher, and A. Singer, "Decoding binary node labels from censored edge measurements: Phase transition and efficient recovery," *Network Science and Engineering, IEEE Transactions on*, vol. 1, no. 1, pp. 10–22, 2014.

[20] B. Hajek, Y. Wu, and J. Xu, "Achieving exact cluster recovery threshold via semidefinite programming: Extensions," arXiv 1502.07738.

[21] A. Saade, M. Lelarge, F. Krzakala, and L. Zdeborová, "Spectral detection in the censored block model," in *Information Theory (ISIT), 2015 IEEE International Symposium on*. IEEE, 2015, pp. 1184–1188.

[22] J. Baik, G. Ben Arous, and S. Péché, "Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices," *Annals of Probability*, pp. 1643–1697, 2005.

[23] R. Latała, "Some estimates of norms of random matrices," *Proceedings of the American Mathematical Society*, vol. 133, no. 5, p. 1273, 2005.

[24] S. Chatterjee, "A generalization of the lindeberg principle," *Ann. Probab.*, vol. 34, no. 6, pp. 2061–2076, 11 2006. [Online]. Available: http://dx.doi.org/10.1214/009117906000000575

[25] S. B. Korada and A. Montanari, "Applications of the lindeberg principle in communications and statistical learning," *IEEE Transactions on Information Theory*, vol. 57, no. 4, pp. 2440–2450, April 2011.

APPENDIX

A. Proof of channel universality

Here we present the detailed proof of channel universality. In this proof, with a slight abuse of notation, we let $\langle \boldsymbol{A}, \boldsymbol{B} \rangle := \sum_{i \le j} A_{ij} B_{ij}$ for two symmetric matrices $\boldsymbol{A}$ and $\boldsymbol{B}$. By definition

$$I(\boldsymbol{W}; \boldsymbol{Y}) = \mathbb{E}_{\boldsymbol{W},\boldsymbol{Y}} \left[\log \frac{P_{\rm out}(\boldsymbol{Y}|\boldsymbol{W})}{\int P(\boldsymbol{W'}) P_{\rm out}(\boldsymbol{Y}|\boldsymbol{W'}) \mathrm{d}\boldsymbol{W'}} \right].$$

Define $D_{ij} = S'_{ij} + 1/\Delta$ and $\mathcal{H}(\boldsymbol{W}, \boldsymbol{Y}) = \langle \boldsymbol{W}, \boldsymbol{S} \rangle - \frac{\|\boldsymbol{W}\|_F^2}{2\Delta} + \frac{1}{2}\langle \boldsymbol{D}, \boldsymbol{W} \circ \boldsymbol{W} \rangle$, where $\circ$ denotes the element-wise matrix product. In view of (9),

$$I(\boldsymbol{W}; \boldsymbol{Y}) = \mathbb{E}_{\boldsymbol{W},\boldsymbol{Y}} \left[\log \frac{e^{\mathcal{H}(\boldsymbol{W},\boldsymbol{Y})}}{\int P(\boldsymbol{W'}) e^{\mathcal{H}(\boldsymbol{W'},\boldsymbol{Y})} \mathrm{d}\boldsymbol{W'}} \right] + O(\sqrt{n}).$$

In the following, we compute $\mathbb{E}_{\boldsymbol{W},\boldsymbol{Y}}\left[\mathcal{H}(\boldsymbol{W}, \boldsymbol{Y})\right]$ and $\mathbb{E}_{\boldsymbol{Y}}\left[\log\left(\int P(\boldsymbol{W'}) e^{\mathcal{H}(\boldsymbol{W'},\boldsymbol{Y})} \mathrm{d}\boldsymbol{W'}\right)\right]$ up to additive errors on the order of $O(\sqrt{n})$.

Lemma A.1.

$$\mathbb{E}_{\boldsymbol{W},\boldsymbol{Y}}\left[\mathcal{H}(\boldsymbol{W}, \boldsymbol{Y})\right] = \frac{n\left(\mathbb{E}[x^2]\right)^2}{4\Delta} + O(\sqrt{n}).$$

Proof: Notice that

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{Y}|\boldsymbol{W}}[\langle \boldsymbol{W}, \boldsymbol{S} \rangle] \\
&= \frac{n(n+1)}{2} W_{12} \mathbb{E}_{Y_{12}|W_{12}}[S_{12}] \\
&= \frac{n(n+1)}{2} W_{12} \mathbb{E}_{Y_{12}|0}[S_{12}(1 + S_{12}W_{12} + O(n^{-1})] \\
&= \frac{n(n+1)}{2} W_{12}^2 \mathbb{E}_{Y_{12}|0}[S_{12}^2] + O(\sqrt{n}) \\
&= \frac{n+1}{2\Delta} x_1^2 x_2^2 + O(\sqrt{n}),
\end{aligned}$$

where we used the fact that $\mathbb{E}_{Y_{12}|0}[S_{12}] = 0$ and $\mathbb{E}_{Y_{12}|0}[S_{12}^2] = 1/\Delta$. It follows that

$$\mathbb{E}_{\boldsymbol{W},\boldsymbol{Y}}[\langle \boldsymbol{W}, \boldsymbol{S} \rangle] = \frac{n+1}{2\Delta}\left(\mathbb{E}[x^2]\right)^2 + O(\sqrt{n}).$$

Also,

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{Y}|\boldsymbol{W}}\left[\langle \boldsymbol{D}, \boldsymbol{W} \circ \boldsymbol{W} \rangle\right] \\
&= \frac{n(n+1)}{2} W_{12}^2 \mathbb{E}_{Y_{12}|W_{12}}[D_{12}] \\
&= \frac{n(n+1)}{2} W_{12}^2 \mathbb{E}_{Y_{12}|0}[D_{12}(1 + O(|S_{12}|n^{-1/2})] \\
&= O(\sqrt{n}),
\end{aligned}$$

where we used the fact that $\mathbb{E}_{Y_{12}|0}[D_{12}] = 0$, and S'_{12} and $\mathbb{E}_{Y_{12}|0}[|S_{12}|]$ are bounded. The lemma readily follows by observing that

$$\mathbb{E}\left[\boldsymbol{W}\|_F^2\right] = \sum_{i \le j} \mathbb{E}[x_i^2 x_j^2] = \frac{n+1}{2\Delta}\left(\mathbb{E}[x^2]\right)^2$$

■

Define

$$\Phi(\boldsymbol{S}) = \log\left(\int P(\boldsymbol{W'}) e^{\langle \boldsymbol{W'}, S\rangle - \frac{\|\boldsymbol{W'}\|_F^2}{2\Delta}} \mathrm{d}\boldsymbol{W'}\right).$$

Lemma A.2.

$$\mathbb{E}_{\boldsymbol{Y}}\left[\log\left(\int P(\boldsymbol{W'}) e^{\mathcal{H}(\boldsymbol{W'},\boldsymbol{Y})} \mathrm{d}\boldsymbol{W'}\right)\right] = \mathbb{E}_{\boldsymbol{Y}}[\Phi(\boldsymbol{S})] + O(\sqrt{n}).$$

Proof: Notice that

$$|\langle \boldsymbol{D}, \boldsymbol{W} \circ \boldsymbol{W} \rangle| = \frac{1}{n}\Big|\sum_{i \le j} D_{ij} x_i^2 x_j^2\Big| \le \frac{1}{n}\|\boldsymbol{D}\|\|\boldsymbol{x}\|_4^4 = O(\|\boldsymbol{D}\|).$$

Thus,

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{Y}}\left[\log\left(\int P(\boldsymbol{W'}) e^{\mathcal{H}(\boldsymbol{W'},\boldsymbol{Y})} \mathrm{d}\boldsymbol{W'}\right)\right] \\
&= \mathbb{E}_{\boldsymbol{Y}}[\Phi(S)] + O\left(\mathbb{E}_{\boldsymbol{Y}}[\|\boldsymbol{D}\|]\right).
\end{aligned}$$

Recall that $D_{ij} = \frac{\partial^2 \log P_{\rm out}(Y_{ij}|0)}{\partial w^2} + \frac{1}{\Delta}$. Then $\boldsymbol{D}$ is a symmetric matrix where $\{D_{ij}\}_{i \le j}$ are independent and identically distributed. Moreover,

$$\mathbb{E}_{\boldsymbol{Y}}[D_{ij}] = \mathbb{E}_{Y_{ij}|0}[D_{ij}(1 + O(|S_{ij}|/\sqrt{n})] = O(1/\sqrt{n}),$$

and since S'_{ij} is bounded, $\mathbb{E}_{\boldsymbol{Y}}[D_{ij}^4] = O(1)$. By Latala's theorem [23], $\mathbb{E}_{\boldsymbol{Y}}[\|\boldsymbol{D} - \mathbb{E}_{\boldsymbol{Y}}[\boldsymbol{D}]\|] = O(\sqrt{n})$. Since $\|\mathbb{E}_{\boldsymbol{Y}}[\boldsymbol{D}]\| = O(\sqrt{n})$. By triangle's inequality, $\mathbb{E}_{\boldsymbol{Y}}[\|\boldsymbol{D}\|] = O(\sqrt{n})$, and the lemma follows. ■

Finally, we apply the generalized Lindeberg principle [24, 25] to show the following lemma. Let

$$\boldsymbol{U} = \frac{\boldsymbol{W}}{\Delta} + \frac{\boldsymbol{\xi}}{\sqrt{\Delta}},$$

where $\xi_{ij} \sim \mathcal{N}(0,1)$ for $i \le j$ and $\xi_{ji} = \xi_{ij}$.

Lemma A.3.

$$\mathbb{E}_{\boldsymbol{Y}}[\Phi(\boldsymbol{S})] = \mathbb{E}_{\boldsymbol{U}}[\Phi(\boldsymbol{U})] + O(\sqrt{n}).$$

Define $a_{ij} = \mathbb{E}[S_{ij}] - \mathbb{E}[U_{ij}]$ and $b_{ij} = \mathbb{E}[S_{ij}^2] - \mathbb{E}[U_{ij}^2]$. Note that

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{Y}|\boldsymbol{W}}[S_{ij}] &= \mathbb{E}_{Y_{ij}|0}[S_{ij}(1 + W_{ij}S_{ij} + O(n^{-1}))] \\
&= \frac{W_{ij}}{\Delta} + O(n^{-1}),
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{Y}|\boldsymbol{W}}[S_{ij}^2] &= \mathbb{E}_{Y_{ij}|0}[S_{ij}^2(1 + O(|S_{ij}|n^{-1/2}))] \\
&= \frac{1}{\Delta} + O(n^{-1/2}),
\end{aligned}$$

where we used the fact that $\mathbb{E}_{Y_{12}|0}[|S_{ij}|^3] = O(1)$. Thus $a_{ij} = O(n^{-1})$ and $b_{ij} = O(n^{-1/2})$. Also, one can check that $\mathbb{E}[|S_{ij}|^3] + \mathbb{E}[|U_{ij}|^3] = O(1)$. Let $[f(\boldsymbol{W})]$ denote the expectation of $f(\boldsymbol{W})$ with respect to the measure defined

by $P(\boldsymbol{W})e^{\mathcal{H}(\boldsymbol{W},\boldsymbol{Y})}\mathrm{d}\boldsymbol{W}/\int P(\boldsymbol{W'})e^{\mathcal{H}(\boldsymbol{W'},\boldsymbol{Y})}\mathrm{d}\boldsymbol{W'}$. It follows that

$$\begin{aligned}
\frac{\partial\Phi}{\partial S_{ij}} &= [W_{ij}] = O(n^{-1/2})\\
\frac{\partial\Phi^2}{\partial S_{ij}^2} &= [W_{ij}^2] - [W_{ij}]^2 = O(n^{-1})\\
\frac{\partial\Phi^3}{\partial S_{ij}^3} &= [W_{ij}^3] - 3[W_{ij}][W_{ij}^2] + 2[W_{ij}]^3 = O(n^{-3/2}).
\end{aligned}$$

Therefore, by Lindeberg principle [24, Theorem 1.1],

$$\begin{aligned}
&\left|\mathbb{E}_{\boldsymbol{Y}}[\Phi(\boldsymbol{S})] - \mathbb{E}_{\boldsymbol{U}}[\Phi(\boldsymbol{U})]\right|\\
&\le O(n^{-1/2})\sum_{i\le j} a_{ij} + O(n^{-1})\sum_{i\le j} b_{ij} + O(\sqrt{n})\\
&= O(\sqrt{n}).
\end{aligned}$$

In conclusion, we have shown that

$$\begin{aligned}
I(\boldsymbol{W};\boldsymbol{Y}) &= \frac{n\left(\mathbb{E}[x^2]\right)^2}{4\Delta} - \mathbb{E}_{\boldsymbol{U}}[\Phi(\boldsymbol{U})] + O(\sqrt{n})\\
&= I(\boldsymbol{W};\boldsymbol{W}+\sqrt{\Delta}\,\boldsymbol{\xi}) + O(\sqrt{n}).
\end{aligned}$$

B. Nishimori identities

A key ingredient in our interpolation proof is the following Nishimori identities [8, 13, 15], which hold for Bayesian inference problems.

Lemma A.4 (Nishimori identities)**.** *Let x^* denote a random sample from a prior distribution $p(x)$, and we observe y randomly generated from $p(y|x^*)$. Let x and x' denote two independent random samples from $p(x|y)$ with the posterior distribution $p(x|y) = p(x)p(y|x)/p(y)$. Then for all f such that $\mathbb{E}[|f(x,x^*)|] < \infty$,*

$$\mathbb{E}[f(x,x')] = \mathbb{E}[f(x,x^*)].$$

Proof: By definition,

$$\begin{aligned}
\mathbb{E}[f(x,x^*)] &= \int p(x^*)\int p(y|x^*)\int f(x,x^*)p(x|y)\mathrm{d}x\mathrm{d}y\mathrm{d}x^*\\
&\overset{(a)}{=} \int p(y)\int\int f(x,x^*)p(x|y)p(x^*|y)\mathrm{d}x\mathrm{d}x^*\mathrm{d}y\\
&= \mathbb{E}[f(x,x')],
\end{aligned}$$

where (a) follows from the fact that $p(x^*)p(y|x^*) = p(x^*|y)p(y)$ and Fubini's theorem. ∎

C. Laplace method

We present the rigorous proof. For fixed $\boldsymbol{x^*}$ and $\boldsymbol{z}$, let

$$G(m,\boldsymbol{x^*},\boldsymbol{z}) = \frac{-m^2}{2\Delta} + \frac{1}{n}\sum_i \mathcal{J}\left(\frac{\widehat{m}}{\Delta}, \frac{mx_i^*}{\Delta} + \sqrt{\frac{\widehat{m}}{\Delta}}z_i\right).$$

and

$$\begin{aligned}
\eta(m,x_i^*,z_i) &= \Delta\times\partial_m\mathcal{J}\left(\frac{\widehat{m}}{\Delta}, \frac{mx_i^*}{\Delta} + \sqrt{\frac{\widehat{m}}{\Delta}}z_i\right)\\
&= [x_i]x_i^*,
\end{aligned}$$

where $[x_i]$ denotes the mean of x_i under the distribution proportional to $\exp\left(-\frac{\widehat{m}x^2}{2\Delta} + x(\frac{mx_i^*}{\Delta} + \sqrt{\frac{\widehat{m}}{\Delta}}z)\right)p(x)d(x)$. Since $p(x)$ has a finite support, η is bounded. Without loss of generality, assume $|\eta| \le C$. It follows that m achieving the maximum value of $G(m,\boldsymbol{x^*},\boldsymbol{z})$ must satisfy

$$m = \frac{1}{n}\sum_i \eta(m,x_i,z_i).$$

Similarly, m achieving the maximum value of $\mathbb{E}\left[G(m,x_1,z_1)\right]$ must satisfy

$$m = \mathbb{E}\left[\eta(m,x_1,z_1)\right].$$

Let $\mathcal{S}$ denote the set of the solutions in $[-C,C]$ of the above fixed point equation. By assumption, $i'_{\mathrm{L}}(m) = 0$ has a finite number of solutions and hence $|\mathcal{S}|$ is finite. Notice that

$$\partial_m\eta(m,x_i,z_i) = \Delta^{-1}(x_i^*)^2\mathsf{var}(x_i) \ge 0.$$

It follows that η is monotone non-decreasing in m. Let $\delta = n^{-1/4}$. Applying [2, Lemma 9], we get that

$$\begin{aligned}
&\mathbb{P}\left\{\sup_{m\in\mathbb{R}}\left|\frac{1}{n}\sum_i \eta(m,x_i,z_i) - \mathbb{E}_{x_1,z_1}[\eta(m,x_1,z_1)]\right| \ge \delta\right\}\\
&\le e^{-\Omega(n^{1/4})}.
\end{aligned}$$

Let $\mathcal{E}$ denote the event that the maximum value of $G(m,x_1,z_1)$ must be attained in the set $\mathcal{G} = \cup_{m\in\mathcal{S}}(m-\delta, m+\delta)$. It follows that $\mathbb{P}\left\{\mathcal{E}^c\right\} \le e^{-\Omega(n^{1/4})}$. Hence,

$$\begin{aligned}
&\frac{1}{n}\mathbb{E}[\log\widetilde{Z}_0]\\
&\le \mathbb{E}[\max_{m\in[-C,C]} G(m,\boldsymbol{x^*},\boldsymbol{z})]\\
&\le \mathbb{E}[\max_{m\in[-C,C]} G(m,\boldsymbol{x^*},\boldsymbol{z})|\mathcal{E}]\mathbb{P}\left\{\mathcal{E}\right\}\\
&\quad + \mathbb{E}[\max_{m\in[-C,C]} G(m,\boldsymbol{x^*},\boldsymbol{z})|\mathcal{E}^c]\mathbb{P}\left\{\mathcal{E}^c\right\}\\
&\le \mathbb{E}[\max_{m\in\mathcal{G}} G(m,\boldsymbol{x^*},\boldsymbol{z})] + \left(O(1) + O\left(\mathbb{E}\left[|z_1|\right]\right)\right)\mathbb{P}\left\{\mathcal{E}^c\right\}.
\end{aligned}$$

Taking the limit $n\to\infty$ on both hand sides of the above displayed equation, we have that

$$\limsup_{n\to\infty}\frac{1}{n}\mathbb{E}[\log\widetilde{Z}_0] \le \limsup_{n\to\infty}\mathbb{E}[\max_{m\in\mathcal{G}} G(m,\boldsymbol{x^*},\boldsymbol{z})].$$

For $m\in\mathcal{G}$, $|G'(m,\boldsymbol{x^*},\boldsymbol{z})| = O(1)$. Thus,

$$\max_{m\in\mathcal{G}} G(m,\boldsymbol{x^*},\boldsymbol{z}) = \max_{m\in\mathcal{S}} G(m,\boldsymbol{x^*},\boldsymbol{z}) + O(\delta).$$

It follows that

$$\limsup_{n\to\infty}\frac{1}{n}\mathbb{E}[\log\widetilde{Z}_0] \le \limsup_{n\to\infty}\mathbb{E}[\max_{m\in\mathcal{S}} G(m,\boldsymbol{x^*},\boldsymbol{z})]. \quad (28)$$

Notice that

$$\begin{aligned}
&\mathbb{E}[\max_{m\in\mathcal{S}} G(m,\boldsymbol{x^*},\boldsymbol{z})]\\
&\le \mathbb{E}\left[\max_{m\in\mathcal{S}}\left|G(m,\boldsymbol{x^*},\boldsymbol{z}) - \mathbb{E}\left[G(m,\boldsymbol{x^*},\boldsymbol{z})\right]\right|\right]\\
&\quad + \max_{m\in\mathcal{S}}\mathbb{E}\left[G(m,\boldsymbol{x^*},\boldsymbol{z})\right]. \quad (29)
\end{aligned}$$

Recall that for any fixed $m \in \mathcal{S}$,

$$\begin{aligned}
&\left|G(m, \boldsymbol{x}^*, \boldsymbol{z}) - \mathbb{E}\left[G(m, \boldsymbol{x}^*, \boldsymbol{z})\right]\right| \\
&= \left| \frac{1}{n} \sum_i \mathcal{J}\left(\frac{\widehat{m}}{\Delta}, \frac{m x_i^*}{\Delta} + \sqrt{\frac{\widehat{m}}{\Delta}} z_i\right) \right. \\
&\quad \left. - \mathbb{E}\left[\mathcal{J}\left(\frac{\widehat{m}}{\Delta}, \frac{m x_1^*}{\Delta} + \sqrt{\frac{\widehat{m}}{\Delta}} z_1\right)\right] \right| \\
&:= T(m).
\end{aligned}$$

Using the fact that $|x_1^*|$ is bounded, one can check that $\mathcal{J}\left(\frac{\widehat{m}}{\Delta}, \frac{m x_1^*}{\Delta} + \sqrt{\frac{\widehat{m}}{\Delta}} z_1\right)$ is sub-Gaussian with $O(1)$ sub-Gaussian norm. Thus by Chernoff's bound, for any fixed $m \in \mathcal{S}$ and $t \geq 0$,

$$\mathbb{P}\left\{T(m) > t\right\} \leq e^{-\Omega(nt^2)}.$$

By a union bound, it follows that with probability at most $|\mathcal{S}| e^{-\Omega(nt^2)}$, $\max_{m \in \mathcal{S} T(m)} > t$. It follows that

$$\begin{aligned}
&\mathbb{E}\left[\max_{m \in \mathcal{S}} T(m)\right] \\
&= \int_0^\infty \mathbb{P}\left\{\max_{m \in \mathcal{S}} T(m) > t\right\} \mathrm{d}t \\
&\leq \int_0^\delta \mathrm{d}t + \int_\delta^\infty \mathbb{P}\left\{\max_{m \in \mathcal{S}} T(m) > t\right\} \mathrm{d}t \\
&\leq \delta + |\mathcal{S}| \int_\delta^\infty e^{-\Omega(nt^2)} = o(1).
\end{aligned}$$

In view of (29), we have that

$$\mathbb{E}[\max_{m \in \mathcal{S}} G(m, \boldsymbol{x}^*, \boldsymbol{z})] = \max_{m \in \mathcal{S}} \mathbb{E}\left[G(m, \boldsymbol{x}^*, \boldsymbol{z})\right] + o(1).$$

Combining the above display with (28), it yields that

$$\begin{aligned}
&\limsup_{n \to \infty} \frac{1}{n} \mathbb{E}[\log \widetilde{Z}_0] \\
&\leq \max_{m \in \mathcal{S}} \mathbb{E}\left[G(m, \boldsymbol{x}^*, \boldsymbol{z})\right] \\
&= \max_m \left\{ \frac{-m^2}{2\Delta} + \mathbb{E}_{x,z} \mathcal{J}\left(\frac{\widehat{m}}{\Delta}, \frac{m x^*}{\Delta} + \sqrt{\frac{\widehat{m}}{\Delta}} z\right) \right\},
\end{aligned}$$

which completes the proof.

D. Interpolating the lower bound

Let us repeat the interpolation strategy of sec. IV-C. We start with eq. (27). Computing the derivative with respect to t, we find that

$$\frac{d}{dt} \mathbb{E}_{\boldsymbol{x}^*, \boldsymbol{z}, \boldsymbol{\xi}}\left[\log \widetilde{Z}_t\right] = \mathbb{E}_{\boldsymbol{x}^*, \boldsymbol{z}, \boldsymbol{\xi}}\left[\sum_{i \leq j} \mathbb{E}_t[\widetilde{\mathcal{A}}_{ij}] + \sum_i \mathbb{E}_t[\widetilde{\mathcal{B}}_i]\right],$$

where

$$\begin{aligned}
\widetilde{\mathcal{A}}_{ij} &= -\frac{x_i^2 x_j^2}{2\Delta n} + \frac{x_i x_j}{2\sqrt{n\Delta t}} \xi_{ij} \\
\widetilde{\mathcal{B}}_i &= \frac{\widehat{m}}{2\Delta} x_i^2 - \frac{1}{2} \sqrt{\frac{\widehat{m}}{\Delta(1-t)}} x_i z_i
\end{aligned}$$

Performing again the integration by part leads to

$$\begin{aligned}
\mathbb{E}_{t, \boldsymbol{\xi}}[\widetilde{\mathcal{A}}_{ij}] &= -\mathbb{E}_{t, \boldsymbol{\xi}}\left[\frac{x_i x_j \mathbb{E}_t[x_i x_j]}{2n\Delta}\right] \\
\mathbb{E}_{t, \boldsymbol{z}}[\widetilde{\mathcal{B}}_i] &= \mathbb{E}_{t, \boldsymbol{z}}\left[x_i \mathbb{E}_t[x_i] \frac{\widehat{m}}{2\Delta}\right].
\end{aligned}$$

Let x' be an independent copy of x. Therefore we have

$$\begin{aligned}
&-\frac{1}{n} \frac{d}{dt} \mathbb{E}_{\xi, \boldsymbol{z}, \boldsymbol{x}^*}\left[\log \widetilde{Z}_t\right] \\
&= \frac{1}{2n^2\Delta} \sum_{i \leq j} \mathbb{E}[x_i x_j x_i' x_j'] - \frac{\widehat{m}}{2\Delta} \sum_i \mathbb{E}[x_i x_i'] \\
&\geq \frac{1}{4n^2\Delta} \mathbb{E}[\langle \boldsymbol{x}, \boldsymbol{x'} \rangle^2] - \frac{\widehat{m}}{2\Delta n} \mathbb{E}[\langle \boldsymbol{x}, \boldsymbol{x'} \rangle] \\
&= \frac{1}{4\Delta} \mathbb{E}\left[\langle \boldsymbol{x}, \boldsymbol{x'} \rangle / n - \widehat{m}\right] - \frac{\widehat{m}^2}{4\Delta} \geq -\frac{\widehat{m}^2}{4\Delta},
\end{aligned}$$

which, with (15) and (26), leads to Theorem (I.3).